The Impact of Internet Services on the Telecommunications Infrastructure

Russ Bell, P.E.
Director, Communications Technology
Advanced Micro Devices, Inc.
russ.bell@amd.com
512-602-5488

Agenda

- Corporate Overview
- Internet Background
- Infrastructure Description
- The World-Wide Wait
- Some Possible Cures
- The Future
- Conclusions

AMD at a Glance

- Fourth Largest US Semiconductor Company
  - Sales 2.4 Billion Dollars
  - Second largest supplier of 32 bit microprocessors
  - 40% of AMD business - communications
  - Number one LAN IC supplier 1995*
  - Largest Line card circuit producer** 150 million shipped to 70 countries
  - CT2 cordless, 802.11 wireless LAN, ISDN, FDDI

* per Dataquest, 1995
** Non-captive

Internet Background

- The communications story of the decade
- A key to the fifth generation of computing
- Huge growth and migration to dial-up network began to reveal cracks in the infrastructure
- Infrastructure
  - legacy protocols, transmission facilities, economic models, etc., not optimal
  - Without the legacy heritage, little would have happened

Internet Architecture

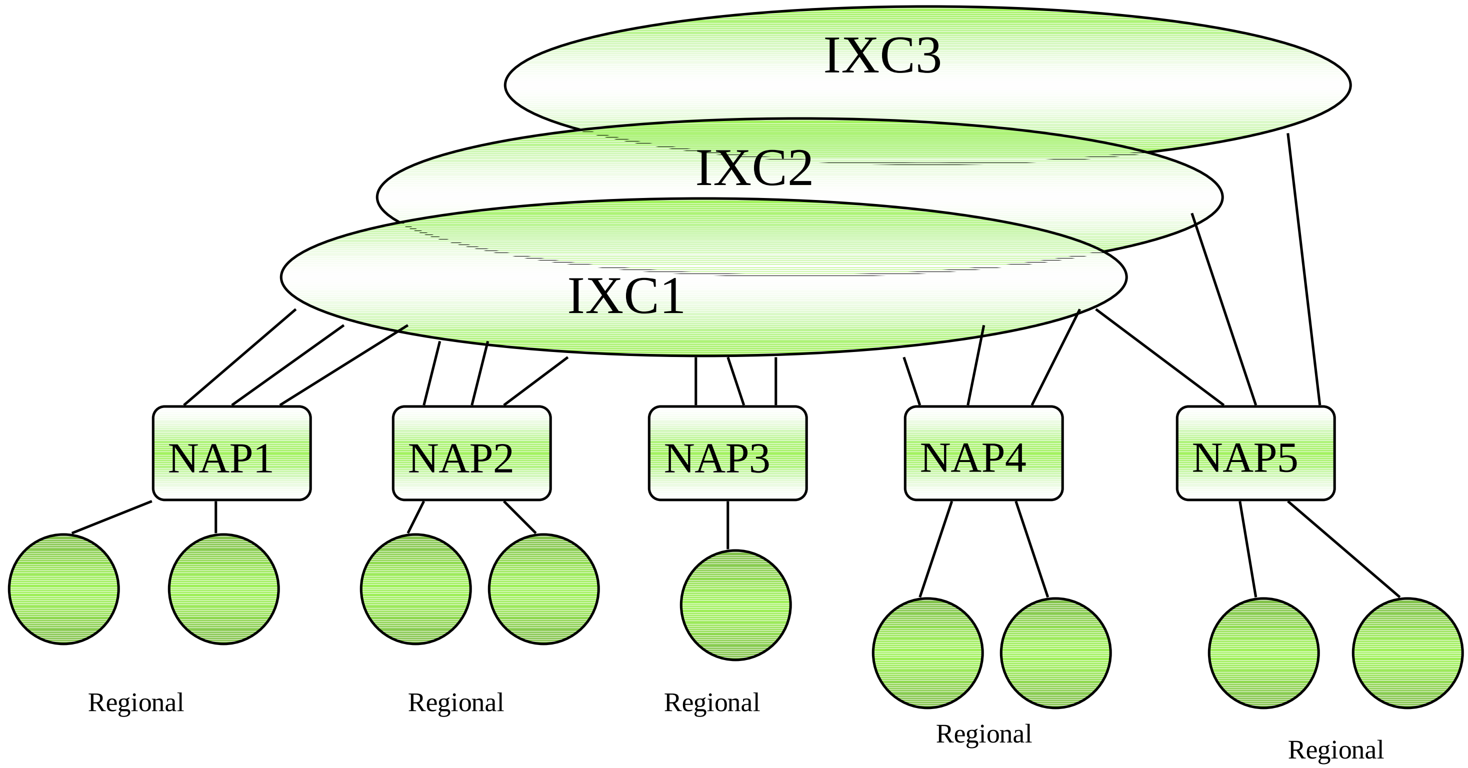

* Regional Networks connect to NAPs
* IXCs connect to all NAPs and provide transit from IXC to IXC

Architectural Issues

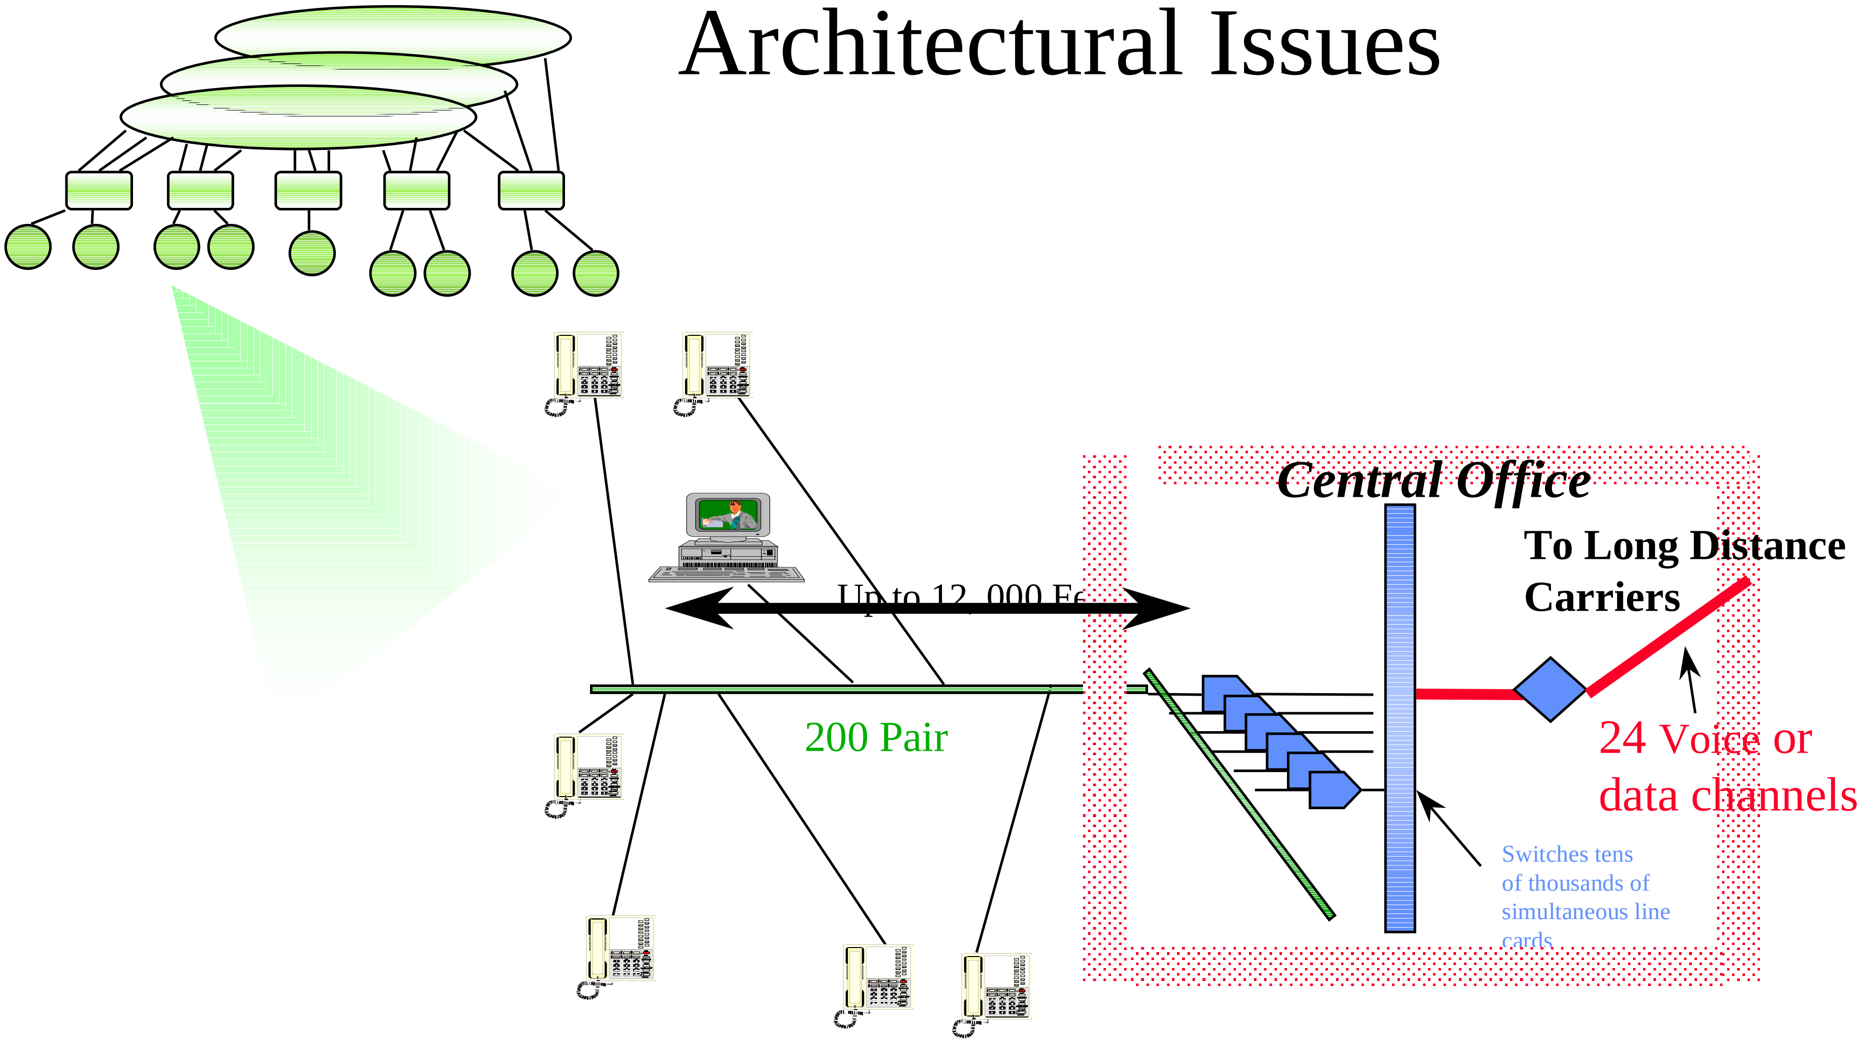

Viewing a WEB page

- Requires the transmission of somewhere near 12 packets over 5 separate networks
- Steps:
  - 1) Computer sends information over MODEM and phone lines to ISP
  - 2) ISP connects to a National backbone provider like MCI, Sprint, BBN
  - 3) The backbone provider passes your packets through one of the peering locations which connects your message to another backbone provider

Viewing a WEB page

- Steps (cont).
  - 4) The second backbone provider connects with the web site's ISP
  - 5) The ISP connects to the appropriate server
- This process is repeated for the return trip

Importance of Response Time

- Georgia Tech Survey Indicates Speed as the number one problem with Internet Service
- Studies Show
  - 3 seconds or less response time maintain user concentration
  - > 10 Seconds results in brain fade

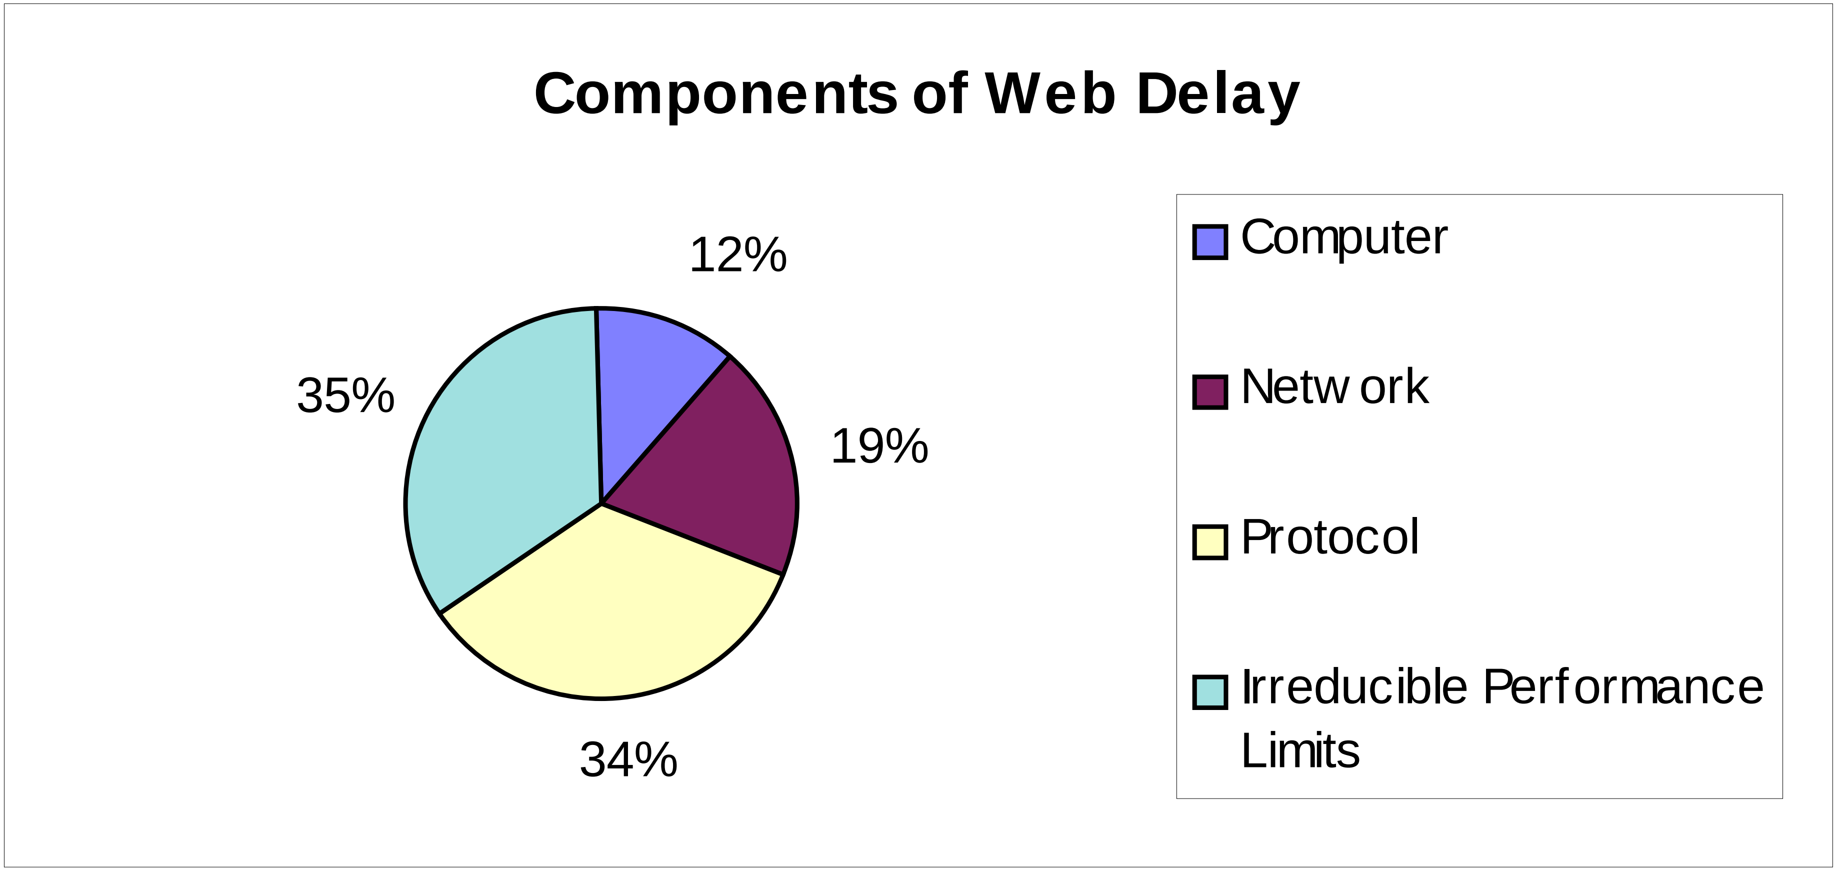

Source: **N.E. Consulting Resources**, Peter Sevcik

Web Response Time

- Depends on
  - Page Complexity (a function of page design)
  - Network Infrastructure
    - Distance Traveled
    - Number of Router Hops
    - Power of Server
    - Protocol Settings (windows, etc.)
    - Access Line Rates

Typical Web Page Delay for Home Users

(Source: BCR 3-96)

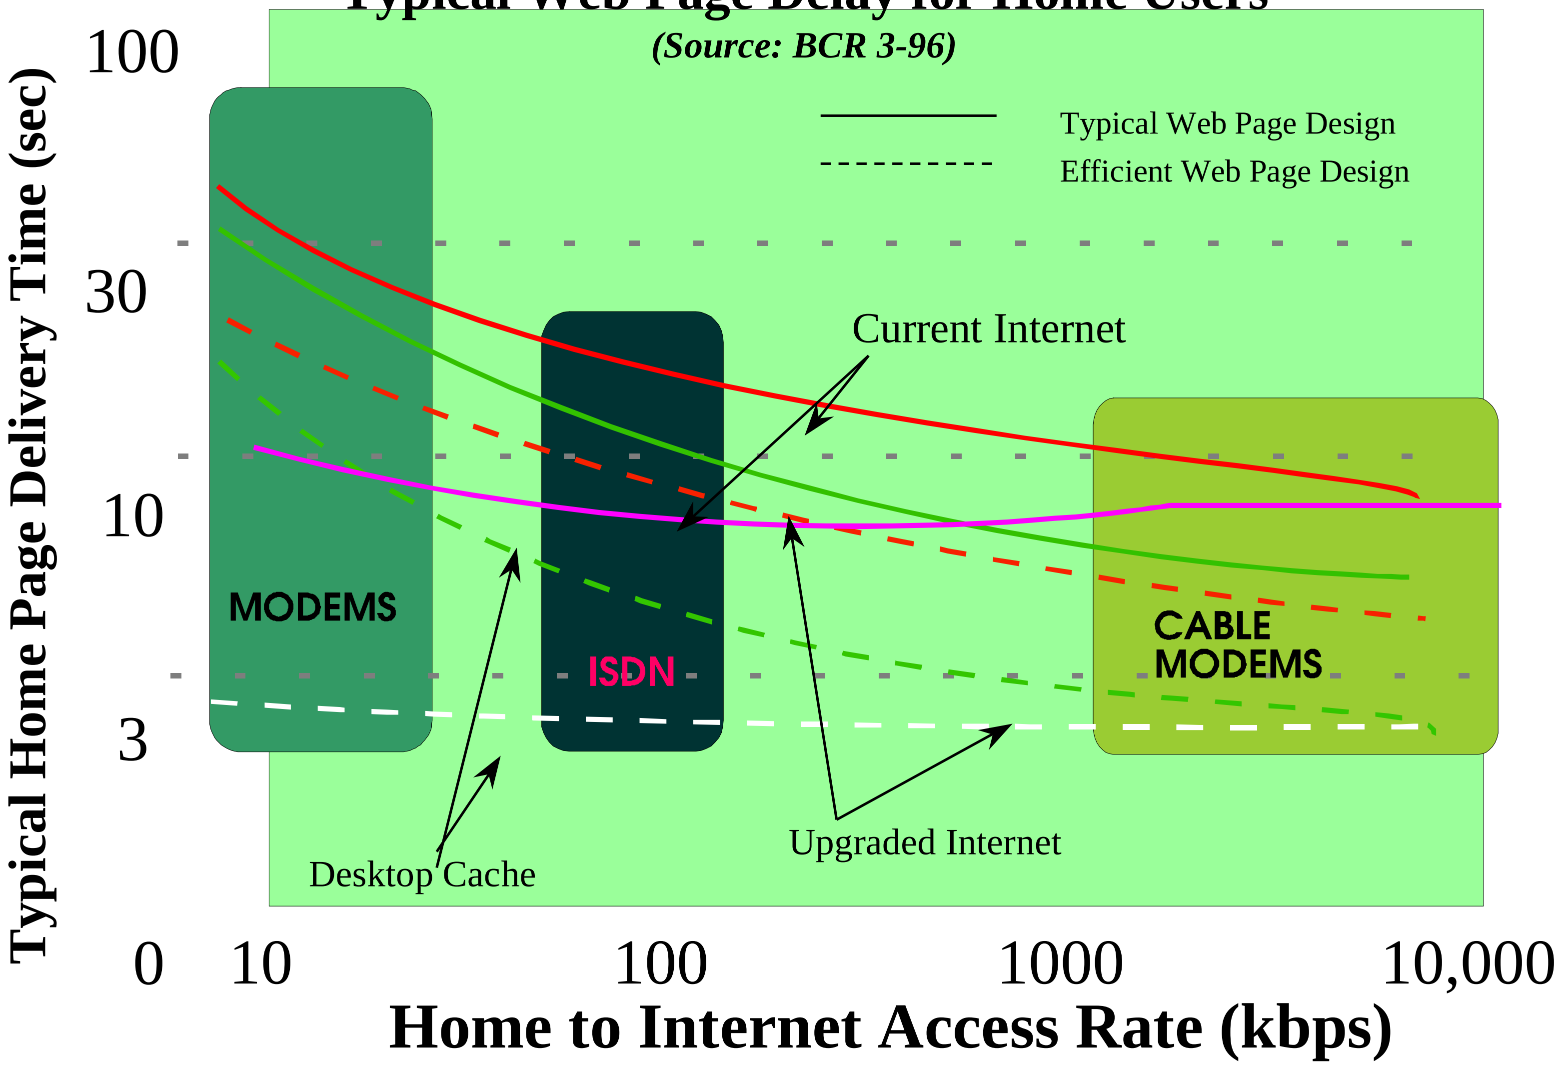

Home Users

- Figure shows time to display basic home pages using various internet access technologies
  - MODEM
  - ISDN
  - Cable Modems
- Typical Home User access contains high graphic content

Home Internet Access

- Solid lines depict typical web transactions
  - Require many overhead transactions to deliver 50k bytes of text and graphics
- Dotted Lines depict - delivery of same 50k bytes with better designed web pages (i.e. those with less overhead)
- In each of these cases, one can deliver the information via
  - Standard Internet
  - Upgraded Internet
  - Desktop Cache

Conclusions

- Over 100 Parameters impact web performance
- ISDN will not deliver in less than 10 seconds
- Cable Modems will still take 12 seconds to deliver a home page
- Remember - this is just for a home page
  - Other streaming applications will probably fare better
  - Local content and caching will be important

Assumptions of “Upgraded Internet”

- Corporate Server is replicated in several regional areas
- Select an ISP with trunking and access rates T3 or greater (45 Mbps)
- Optimized Protocols
- Provides significant performance improvement but is expensive

Desktop Cache

- Performance increase achieved by:
  - Configuring navigator to cache locally
  - Curves assume standard internet access
  - 10 Mbytes of local storage needed for successful operation
  - Local Caching provides excellent performance

Caching and Network Appliances

- Caching analysis says that “Networking Computers” for internet access will not provide optimal performance if they do not contain a hard drive or non-volatile storage for desktop caching.
- These “browsers” have the potential of being slower than traditional desktop PCs capable of supporting caching

Protocol Inefficiencies

- HTTP - very inefficient
- Client Receives HTML Document via HTTP request
  - discovers various links inside HTML document
  - Issues separate request for each embedded HTTP
  - HTML page has been measured at 6.4 kbytes + 7 images per page *

* - Cunha see reference

Net Catches

- TCP IP Protocol impacts performance
  - Silly Window Syndrome
  - TIME_WAIT parameter when connection is closed delays reuse of system resources for 4 minutes
    - TCBs - transmission control buffers
    - Sockets
  - Asymmetrical Services and ACKs

More Net Catches

- TCP maintains state information inside of TCB (transmission control block) for several minutes after the connection has been closed
  - required to prevent data corruption due to sequence number reuse
  - results in thousands of connection records for busy servers and resource limitations “host contacted - waiting response”
  - these TCBs can range in size from 8-32 kbytes
  - a web page with 9 ICONs will have 10 TCBs
  - the transfer may be completed in 1 minute- TCBs remain “tied up” for over 4 minutes

Net Catches (cont.)

- Poorly designed pages can result in a host that runs out of TCBs
- Results in the message “ Host contacted, waiting reply”
- Major cause of user’s perception of a slow internet performance
- What helps?
  – reduce the number of objects on a page (less TCBs)
  – raise the number of TCBs
  – UNIX OS tends to be better at this

Fixes

- Add Main Memory in Servers
- Allocation Table “Tweak” in OS
- HTTP ver 1.1 suggests TCP (T/TCP) - Transactional TCP
  - Allows TCP connection to stay open for several HTTP transfers

When is 100 Mbps not 100 Mbps?

- Purdue researchers D. Comer and J. Lin perform FTP using ATM
  - FTP over TCP IP using ATM transport
  - Standard ATM switch SVC and PVC
  - 100 Mbps fiber optic link between servers and switch
  - ftp://ftp.cs.purdue.edu/pub/comer/TCP.over.ATM.ps.Z

Results

- FTP configured with large buffer size (51k)
- No other applications running
- Measured results of 469 kbps (less than 0.47% of the link throughput)
- Replaced the ATM link with Ethernet and results were 3.8 Mbps
- ATM is 12% as fast as Ethernet?

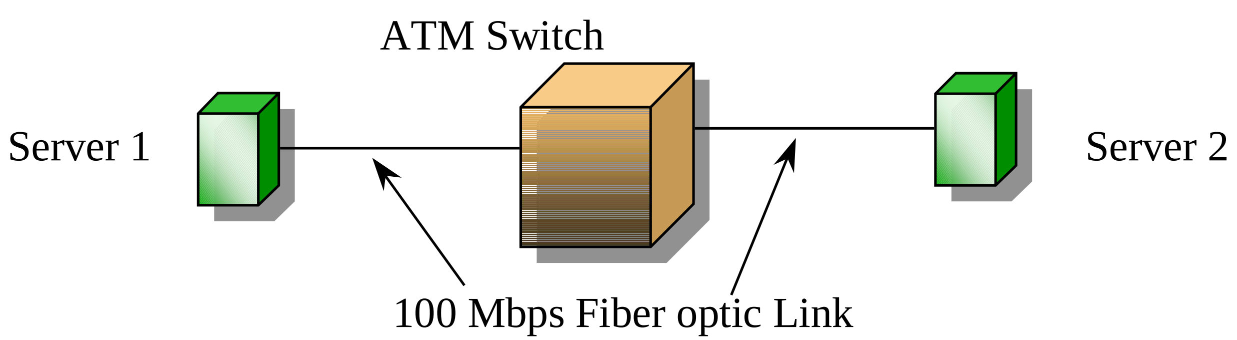

100 not equal to 100?

- Complex interaction
  - ATM buffer size interacts with the SWS of TCP IP
  - One packet is sent based on transmit buffer size relative to OS
  - Awaits ACK from remote
  - Remote SWS Heuristic is tuned to wait 200 ms
  - Link is idle for 200 ms after transmission of 8 k bytes (0.64 ms) for a duty cycle of 0.3%

100 not equal to 100?

- Remember
  - FTP and other TCP dependent protocols tuned to operate over other links
  - Higher speeds, alternate buffering or segmentation strategies, and asymmetrical data rates all impact performance
  - Do not assume legacy TCP applications can be easily upgraded to operate over new transports

Possible Cures

- Cures?
  - Legacy systems have their own unique issues
    - marketing
    - economics
    - technology
    - existing infrastructure
    - regulatory issues
    - business models
    - service reputation

Solutions - Cable

- Loop Solutions for internet access
  - Cable
    - high debt
    - non-standard transmission technology - economies of scale
    - poor service reputation
    - high cost for upgrades

Solutions - ISDN

- The phone company's initial entry
- Ordering information and software configuration
- Initially non-standard implementations
- Islands and signaling were a problem
- Extended reach into the sweet spot of the market
- May exacerbate blocking problem

Solutions - ADSL

- Provides high speed access using existing infrastructure
- Requires loop modifications
  - loading coils and splitters at CO and CPE
- Operates at asymmetrical data rates and may require tuning TCP stack for optimal performance
- Removes long call holding times from CO switch
- Operates simultaneously with voice
- May be initial entree of competitors into the unbundled market

DLCs and AIN

- Digital Loop Carriers Use Advanced Intelligent Network
  - Intercept the callers destination number
  - Determine if it requires excessive holding time
  - Reroute traffic around CO switch
  - Additional capital outlay for RBOC - questionable ROI

Wireless

- LMDS and MMDS offer potential of high speed wireless local loop access
- Rapid deployment
- Unproved Technology and cost points
- Jury is out
- Mobile Access via PCS bands doubtful
- Fixed Access via bandwidth aggregation possible

In the Future

- Network remains slow until infrastructure catches up with demand
- Multimedia Applications are painful and likely to remain so for the near future
- Non - TCP protocols or TCP extensions allow incremental improvements
  - RTP (real time protocol) allows adaptive feedback and tuning of data rates and buffer size
  - RSVP the IETF requires router modifications, authentication and usage based billing - Not likely to be deployed for years

In the Future

- TCP and IP over ATM
  - Experimentation with traffic types, flow detection, QOS, and transport continue and begin to segment the user base
- Alliances of strange bedfellows
  - Cisco / Intel and MCI
- Erosion of Telephony Margins
  - Internet transport of various data type, voice, FAX

Conclusions

- Rapid growth of the net has highlighted the flaws in existing infrastructure
- Beware of "quick fixes"
- Huge momentum of user base, infrastructure suppliers, operators, etc.
- Opportunity for radical change is at hand
  - the net has changed the communications landscape forever
- Patience and savvy are key

References

References

Schulzinne, Henning. "World Wide Web: Whence Whither, What Next?" IEEE Network, March, April 1996, 10-17

C. Cuhna, A. Bestavros, and M. Crovella, "Characteristics of WWW Client Based Traces," Tech. Pub Bu-CS-95-010, Boston University, Computer Science Department, July 1995

Sevcik, Peter. "Designing a high-performance web site," Business Communications Review, March 1996,27-31

D. Comer, J. Lin. "TCP Buffering and Performance over an ATM network," Journal of Internetworking: Research and Experience, vol. 6(1), March 1995 pp1-13

Moskowitz, Robert G. "Why in the world is the web so slow?" Network Computing March 15, 1996, 22-24

Stevens, Richard. TCP IP Illustrated, Volume 1 New York: Addison Wesley Publishing, 1994

Alphanet - *Internet Telephony* Architecture

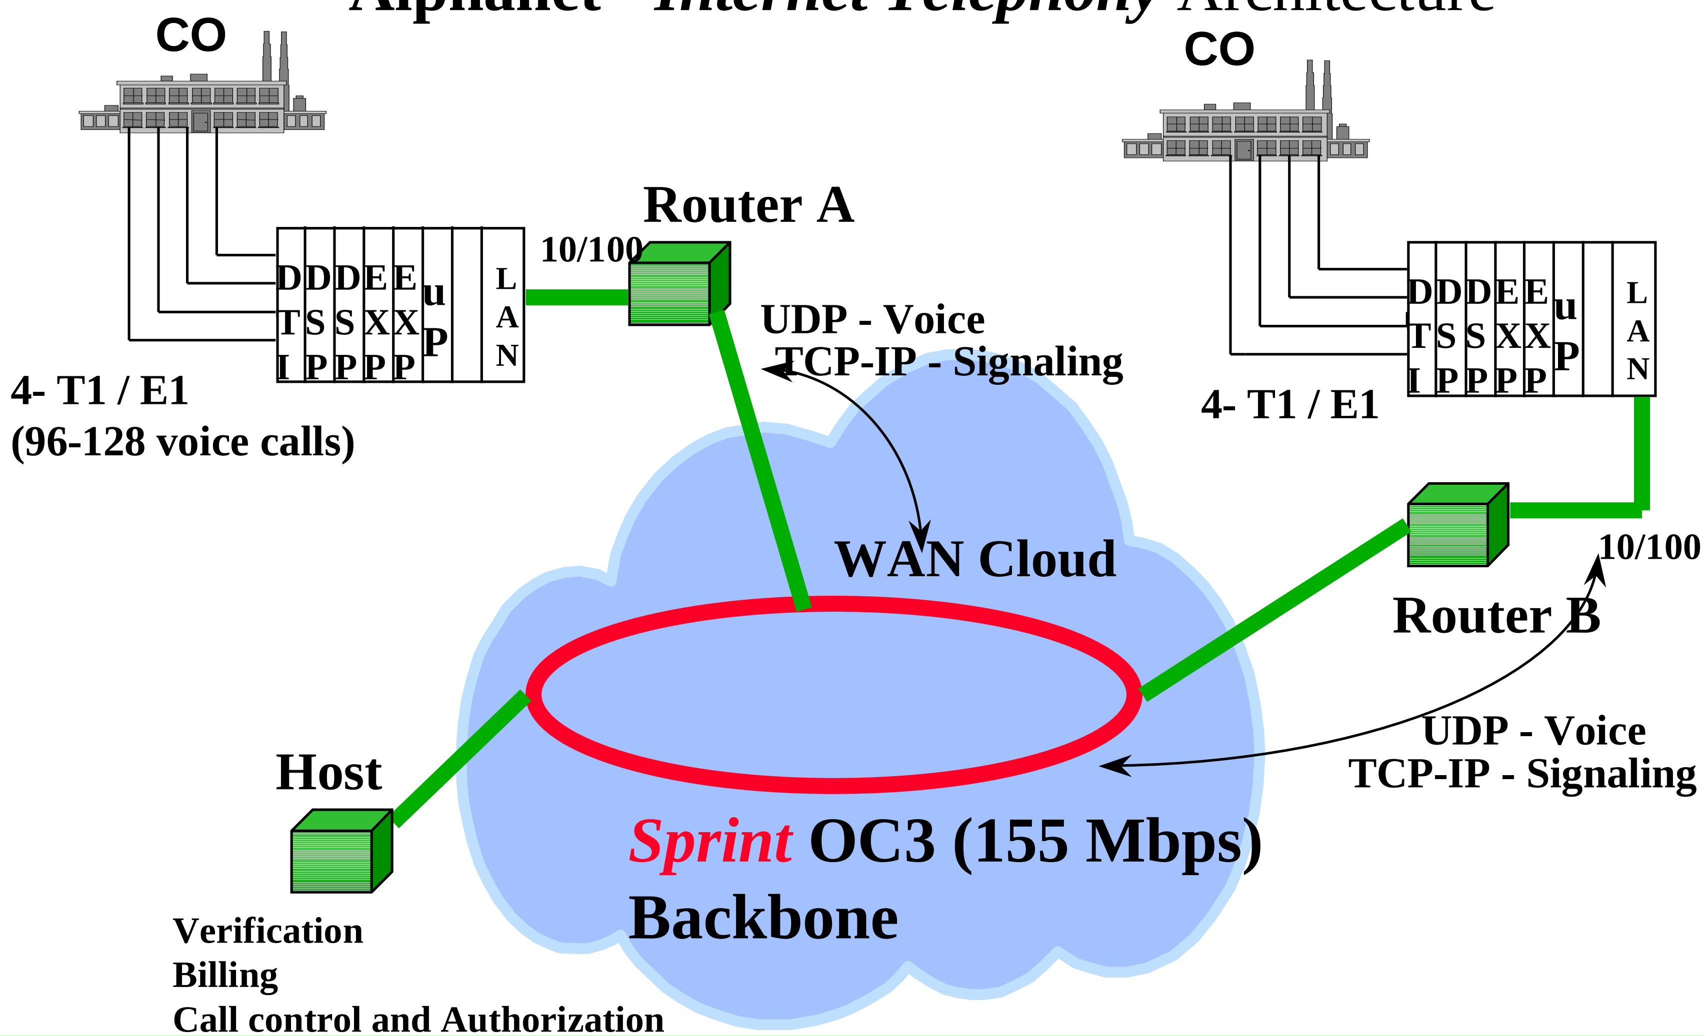

xDSL

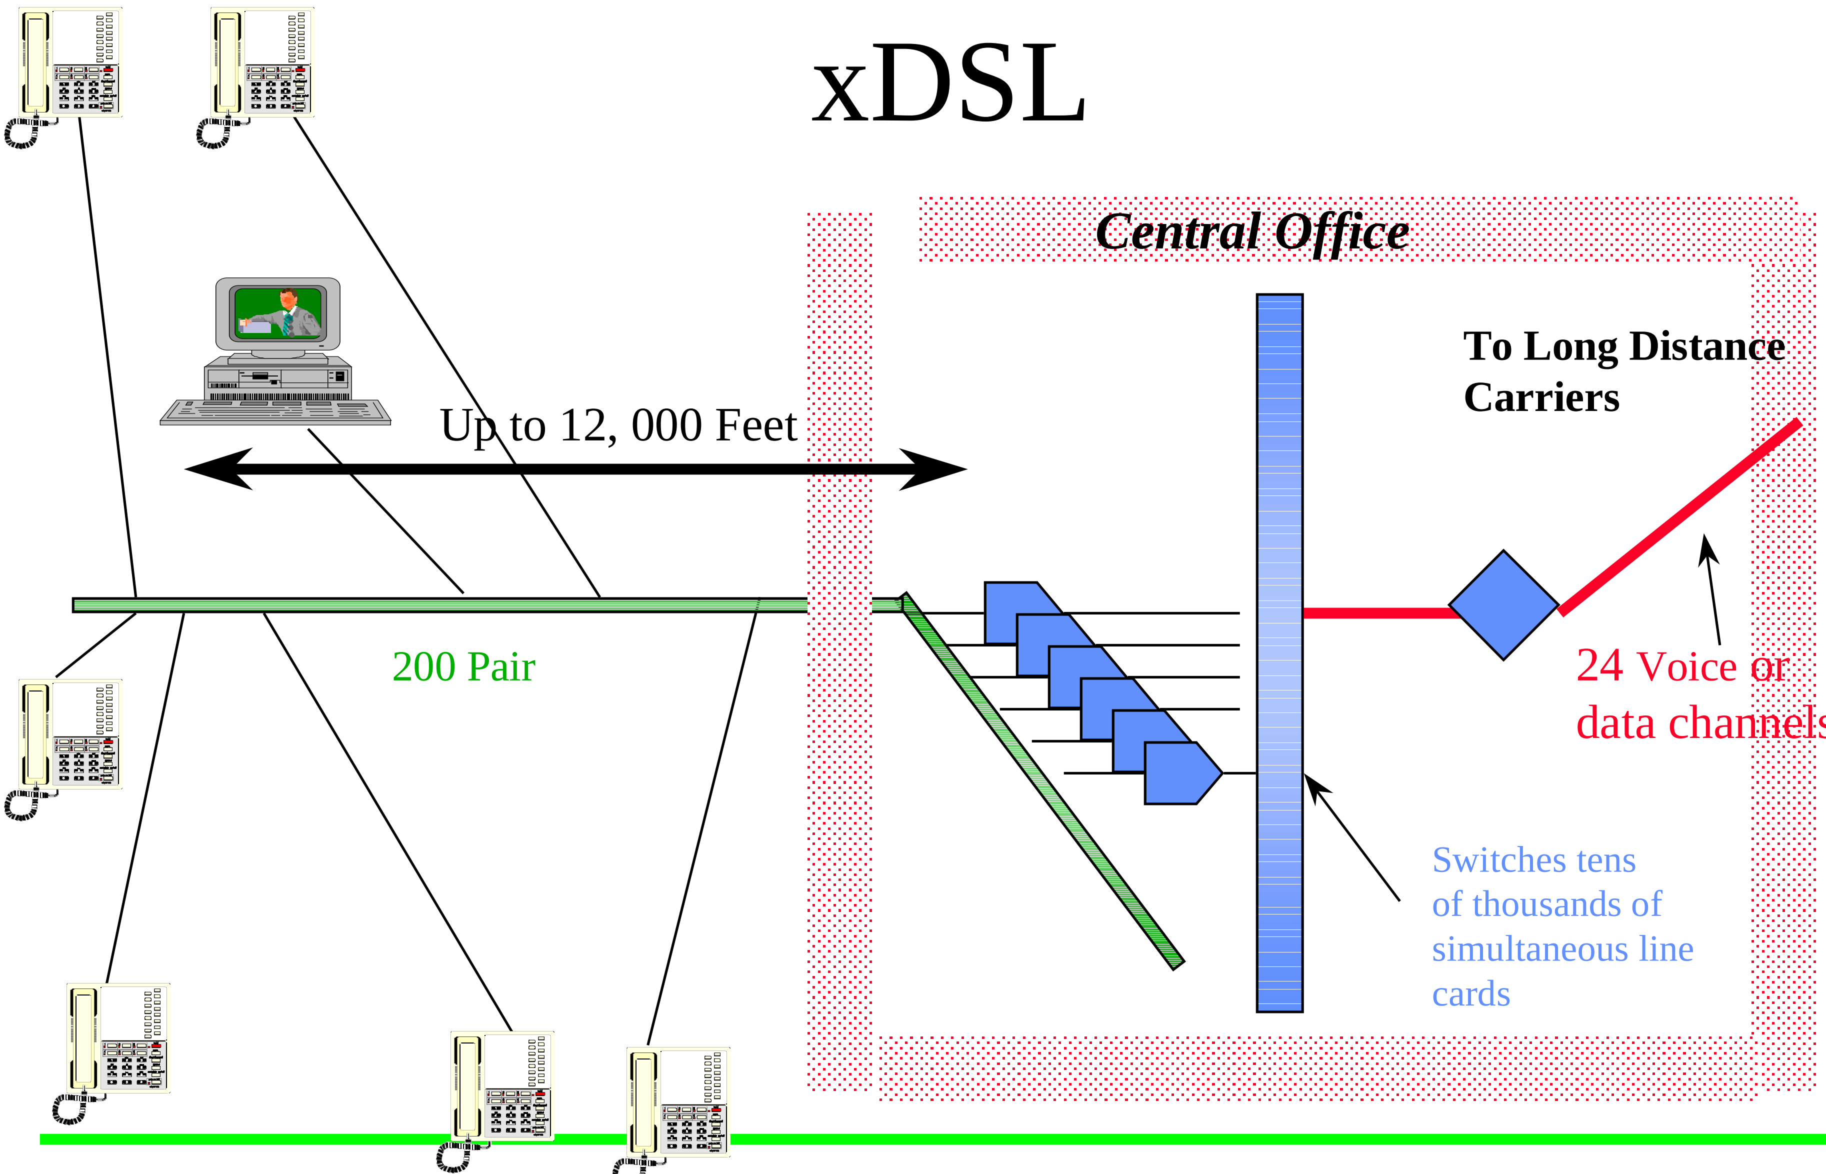

Central Office
To Long Distance Carriers
Up to 12, 000 Feet
200 Pair
24 Voice or data channels
Switches tens of thousands of simultaneous line cards

Unbundling

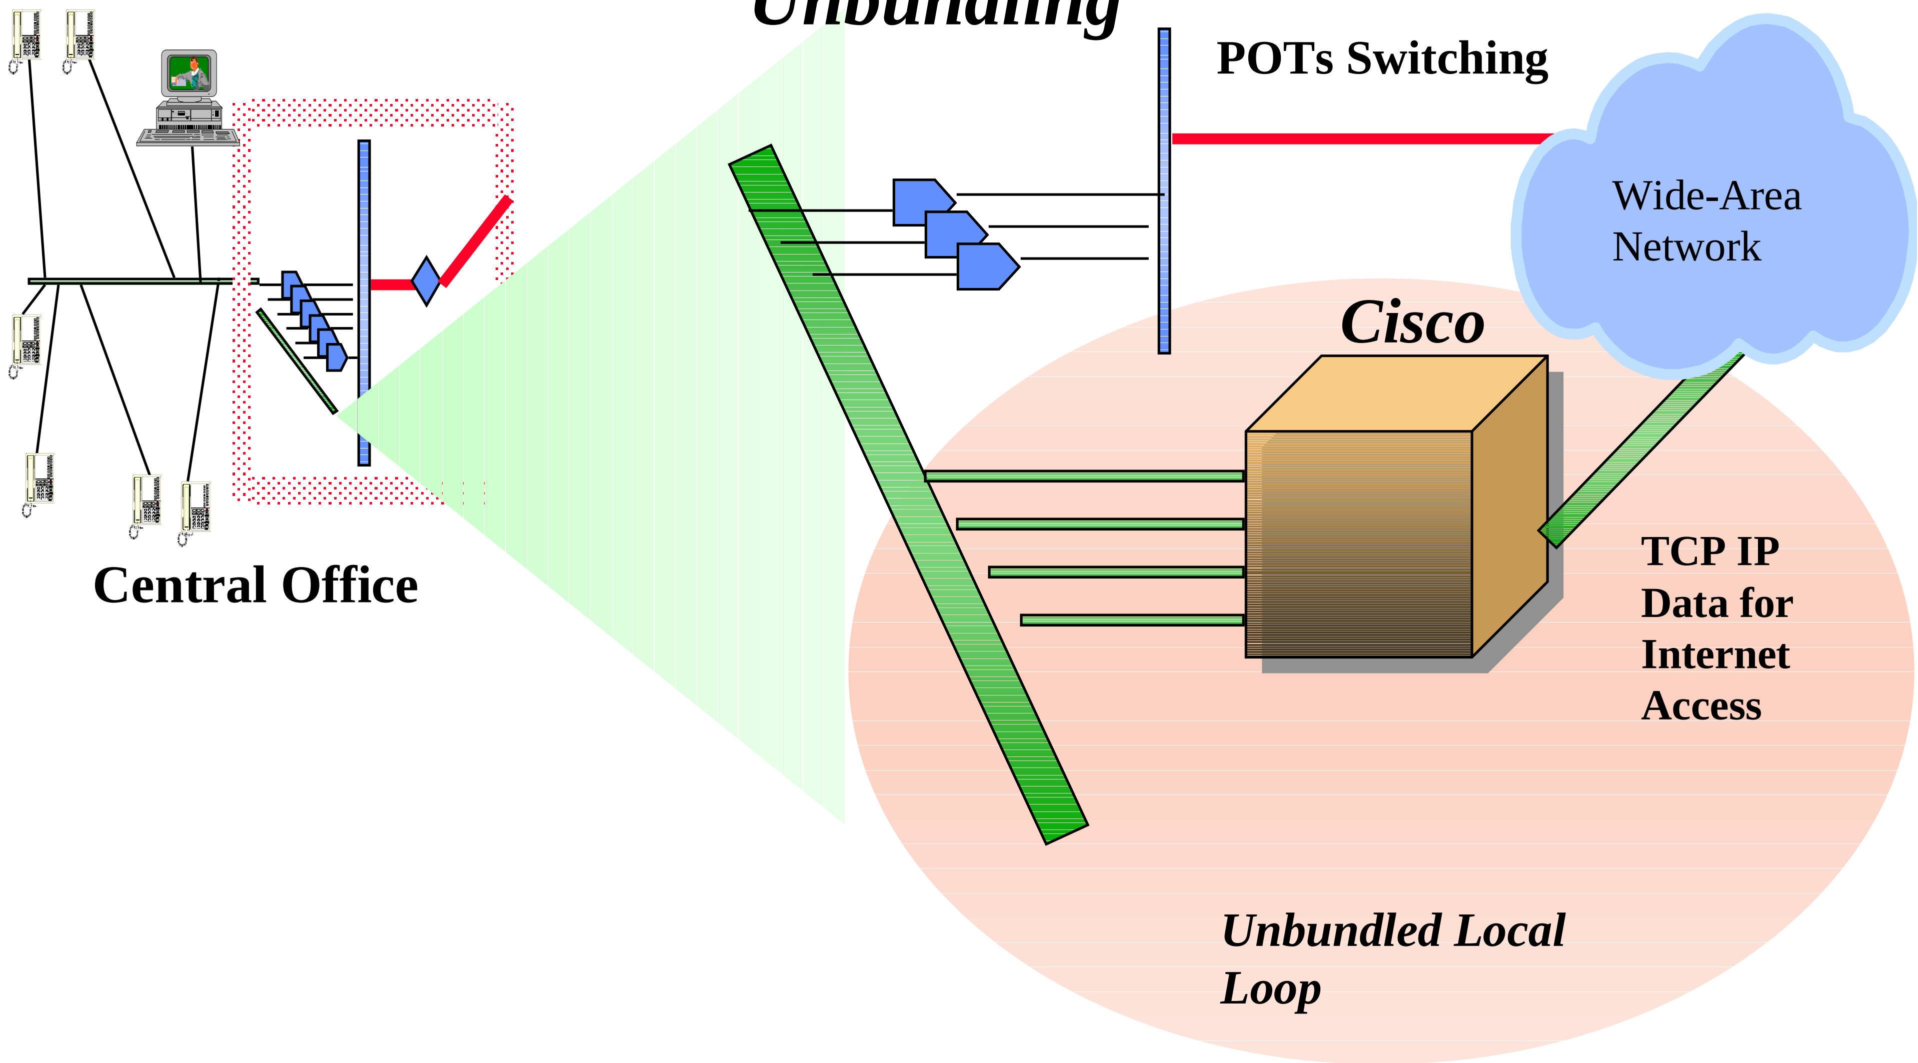
POTs Switching
Wide-Area
Network
Cisco
TCP IP
Data for
Internet
Access
Central Office
Unbundled Local
Loop

Removing the Bottleneck

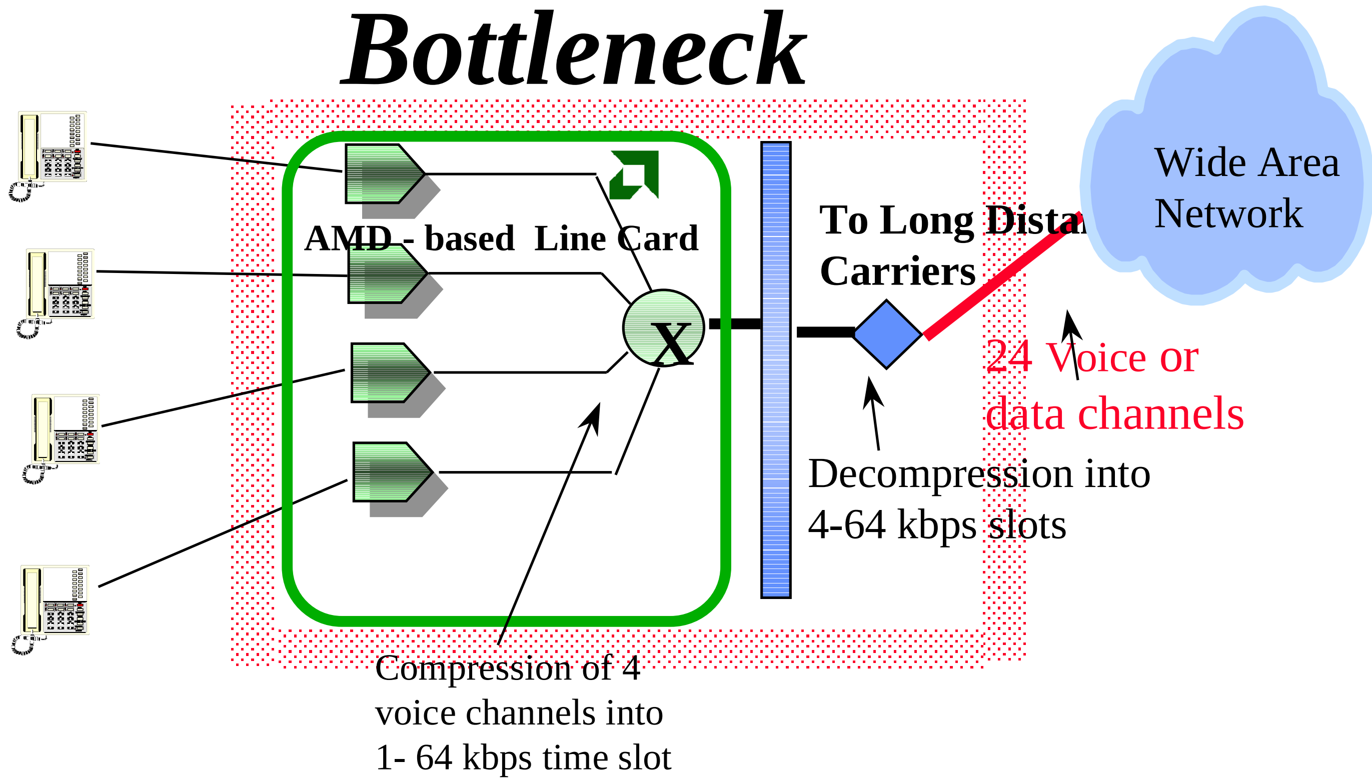

Work with Bellcore and RBOCs to Define Next Generation Platform that preserves “stranded” cost in the network

IP "Switching"

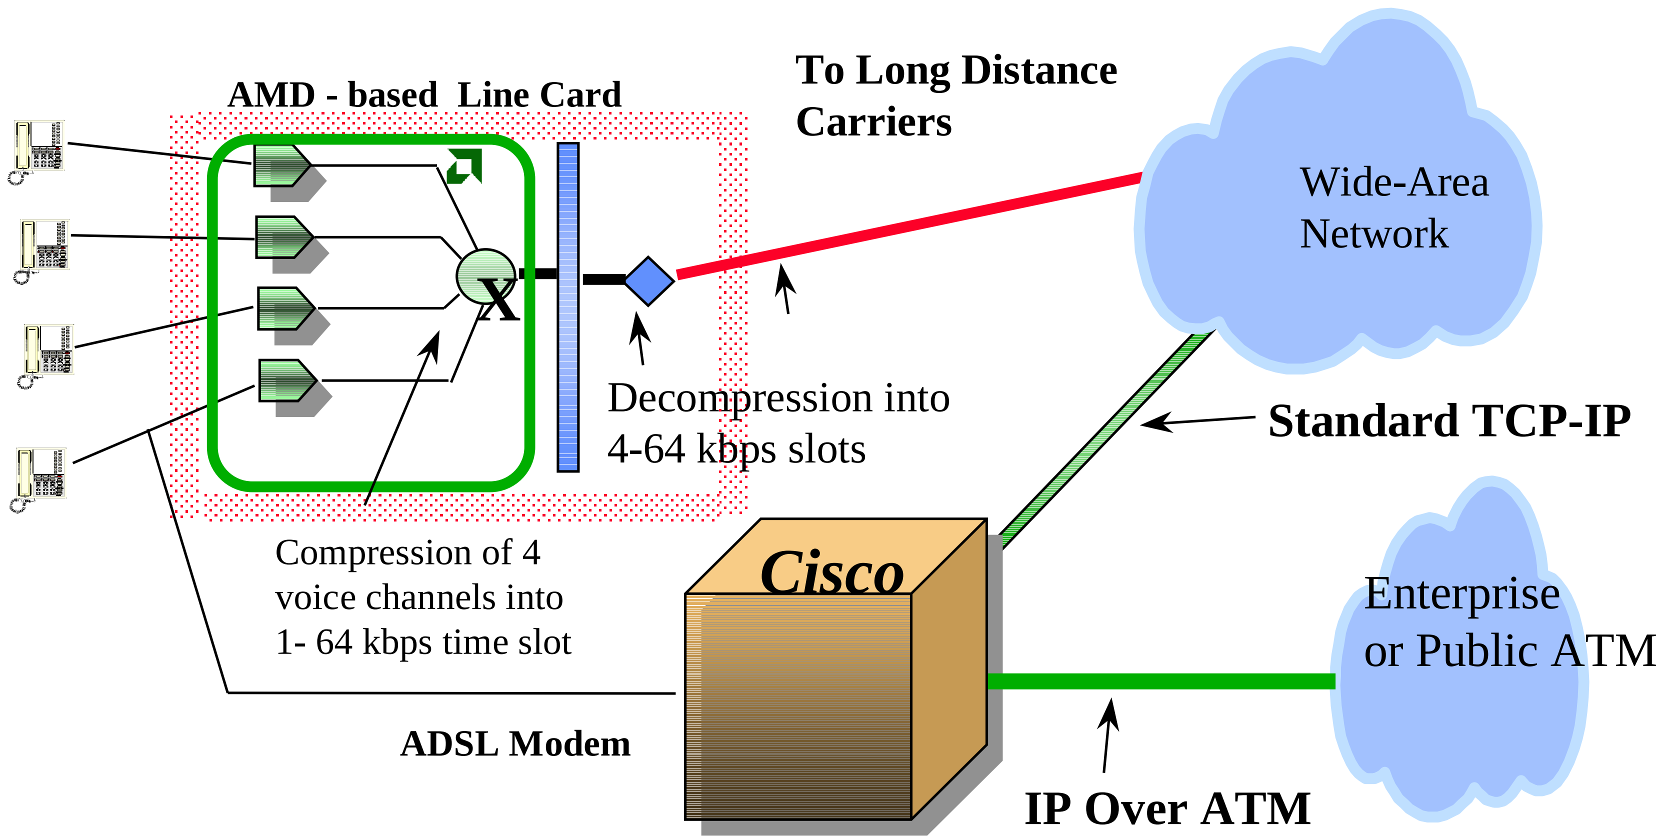

Router Flow Switching Statistics*

- Cisco 7505 Router Between St. Louis and Kansas City
  - A Flow is one session
  - 39 Hours = 58 Million flows, 55% were from HTTP
  - Average Web Flow was 41 seconds long
  - 68% was idle time

* Network Computing - March 15, 1996

Router Analysis

- Average Web Flow - 13 packets / 3563 bytes
- 750 Bytes of Overhead
- Packet flow is insufficient for Van Jacobsen "slow start" to have any effect
- VJSS - sends TCP packets slowly and speeds up after several packets have ACKed. It then speeds up until the ACKS become slower (indicating network congestion) at which time it slows down.

Router Analysis

- Small Packet flows result in congestion and loss in the core routers
- Results in the retransmission of many packets
- Long running sessions like FTP and HTTP transfers slow down as they suffer from the congestion of of short HTTP sessions

Net Math Cont.

- TCP Congestion Control Enters Picture
  - reduces throughput until TCP window is fully opened without regard for link speed
  - Experiments show:
    - Using TCP to transfer 2k byte packets with 70 ms RT delay results in throughput of 10% of max
    - This increases to 50% when 20k byte packets are used

Net Math

- Each HTTP Retrieval Requires
  - At least one round trip time +
  - TCP set up time of 3 round trip delays
  - TCP Connection termination of 2 round trip delays
  - Overlap is possible - smallest with overlap is 4 RT delays
  - Assuming 300 ms RT delay, minimum latency with infinitely fast link = 1.2 s

WWW and its impact on the Internet

- WWW Traffic accounts for over 50% of the traffic on the internet
- Stresses Internet performance as humans are waiting for responses
- Data can be anything from a few bytes to a few megabytes with varying QOS requirements
- There is poor spatial correlation on hyperlinks that are derived from an HTTP retrieval

100 Mbps FTP Yields?

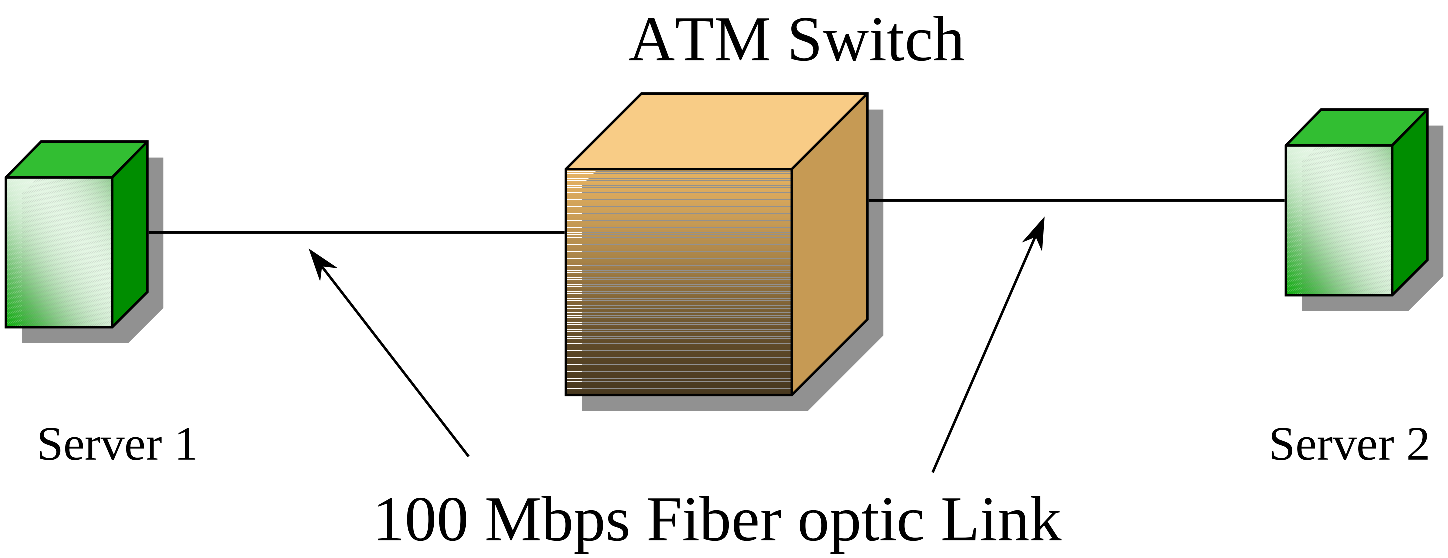

Internet Impact of WWW

- Low Spatial Correlation will strain any IP-over ATM mechanism that attempts to set up individual switched virtual circuits for each web retrieval *
- The net can only scale if the information content is mirrored or cached
  - Mirrors are complete copies that are trusted and updated by the master copy
  - Caches placed between the WWW server and the client